

ORIGINAL

PATENT APPLICATION

ATTORNEY DOCKET NO. 10991147-1

IN THE
UNITED STATES PATENT AND TRADEMARK OFFICE



Inventor(s): Qiming CHEN et al.

Confirmation No.: 5742

Application No.: 09/524,140

Examiner: E. T. Shaffer

Filing Date: 03/10/2000

Group Art Unit: 3623

Title: DISTRIBUTED OLAP-BASED ASSOCIATION RULE GENERATION METHOD AND SYSTEM

Mail Stop Appeal Brief-Patents
Commissioner For Patents
PO Box 1450
Alexandria, VA 22313-1450

TRANSMITTAL OF APPEAL BRIEF

Sir:

Transmitted herewith in **triplicate** is the Appeal Brief in this application with respect to the Notice of Appeal filed on 04/19/2004.

The fee for filing this Appeal Brief is (37 CFR 1.17(c)) \$330.00.

(complete (a) or (b) as applicable)

The proceedings herein are for a patent application and the provisions of 37 CFR 1.136(a) apply.

() (a) Applicant petitions for an extension of time under 37 CFR 1.136 (fees: 37 CFR 1.17(a)-(d) for the total number of months checked below:

() one month	\$110.00
() two months	\$420.00
() three months	\$950.00
() four months	\$1480.00

() The extension fee has already been filled in this application.

(X) (b) Applicant believes that no extension of time is required. However, this conditional petition is being made to provide for the possibility that applicant has inadvertently overlooked the need for a petition and fee for extension of time.

Please charge to Deposit Account **08-2025** the sum of \$330.00. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16 through 1.21 inclusive, and any other sections in Title 37 of the Code of Federal Regulations that may regulate fees. A duplicate copy of this sheet is enclosed.

(X) I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to:
Commissioner for Patents, Alexandria, VA
22313-1450. Date of Deposit: 06/11/2004
OR

() I hereby certify that this paper is being transmitted to the Patent and Trademark Office facsimile number _____ on _____

Number of pages:

Typed Name: Christina L. Paz

Signature: Christina L. Paz

Respectfully submitted,

Qiming CHEN et al.

By

Jonathan M. Harris
Jonathan M. Harris

Attorney/Agent for Applicant(s)
Reg. No. **44,144**

Date: **06/11/2004**

Telephone No.: (713) 238-8000



ORIGINAL

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Appellants:	Qiming CHEN et al.	§	Confirmation No.:	5742
Serial No.:	09/524,140	§	Group Art Unit:	3623
Filed:	03/10/2000	§	Examiner:	E.T. Shaffer
For:	Distributed OLAP-Based	§	Docket No.:	10991147-1
	Association Rule	§		
	Generation Method And	§		
	System	§		

APPEAL BRIEF

Mail Stop Appeal Brief – Patents
Commissioner for Patents
PO Box 1450
Alexandria, VA 22313-1450

Date: June 11, 2004

Sir:

Appellants hereby submit this Appeal Brief in connection with the above-identified application. A Notice of Appeal was filed on April 19, 2004.

I. REAL PARTY IN INTEREST

The real party in interest is the Hewlett-Packard Company.

II. RELATED APPEALS AND INTERFERENCES

Appellants are unaware of any related appeals or interferences.

III. STATUS OF THE CLAIMS

Originally filed claims: 1-25.

Claim cancellations: None.

Added claims: 26-29.

Presently pending claims: 1-29.

Presently appealed claims: 1-29.

IV. STATUS OF THE AMENDMENTS

No claims were amended after the Final Office Action dated February 19, 2004.

V. SUMMARY

Data mining refers to the process of discovering trends in usually large amounts of data. Using a combination of machine learning, statistical analysis, modeling techniques, and database technology, data mining may be used to locate patterns and relationships in the data. Association rules between elements of data may be inferred from the patterns and relationships, allowing trends in the data to be more readily appreciated. See page 2, line 25 - page 3, line 3.

Some data mining techniques focus on the analysis of historical data rather than the analysis of continuous, real-time information. See page 3, lines 17-26. Various problems exist with such techniques. For example, the techniques may be incapable of detecting changes in one or more association rules. See page 3, lines 26-32. Further, such techniques may be incapable of processing the amount of data required to deliver real-time, continuous data mining results. Thus, a trend may not be detected until its relevance has diminished. See page 4, line 7-14.

Appellants address these issues by providing a processing system that comprises "a minimum of two layers of data warehouse/OLAP stations: LDOSs and a GDOS...The LDOSs are responsible for local data mining and summarization, while the GDOS is responsible for merging and mining the input data from the LDOSs, and for providing mining results to LDOSs for business applications." Page 7, lines 11-20 and Figure 1. The processing system may be used to generate scoped association rules. See Page 17, line 10 - page 18, line 20. For exemplary purposes, consider the scoped associated $A \rightarrow B$ [scope = customer, region = CA, timeframe = two months]. This exemplary scoped association rule represents the trend that a customer who purchased item A also purchased item B in the state of California within the past two months. Scoped association rules are defined over a particular population or set of populations, thereby facilitating the cooperative mining of association rules between the GDOS and the LDOSs. See page 17, line 10 - page 18, line 20.

VI. ISSUE(S)

The issue in this Appeal is whether claims 1-29 are unpatentable over Megiddo (U.S. Pat. No. 6,182,070) in view of Castelli (U.S. Pat. No. 5,978,788) under 35 U.S.C. § 103(a).

VII. GROUPING OF CLAIMS

The patentability of the independent claims 1, 8, 17, and 26 will be argued separately below, and thus claims 1, 8, 17, and 26 do not stand or fall together.

VIII. ARGUMENT

The Examiner erred in rejecting claims 1, 8, 17, and 26 as being unpatentable over Megiddo in view of Castelli. Both references are summarized below.

A. The Megiddo Reference

Megiddo is directed to a system and method for determining the significance of association rules. See Megiddo at col. 3 lines 1-7. In Megiddo, one or more synthetic databases are generated from a dataset. Each synthetic database contains a plurality of transactions, each transaction contains one or more items. Association rules having a similar probability threshold value may be discovered for each synthetic database. See Megiddo at col. 3, lines 41-51. Thus, the association rule $S \rightarrow T$ may represent the trend that a customer who purchased item S also purchased item T in the same transaction. See Megiddo at col. 8, lines 54-66.

B. The Castelli Reference

Castelli is directed to an apparatus and method for approximating the data stored in a database by generating multiple projections and/or representations from the database. See Castelli at col. 3 lines 1-7. In Castelli, a database is converted into an initial data cube, the initial data cube is split into multiple resolutions, and a reconstructed data cube is ultimately formed. See Castelli, col. 3, lines 8-29.

C. Claim 1

Claim 1 requires, among other features, "generating scoped association cubes, a population cube and a base cube based on the volume cube, wherein the scoped association cubes comprise a plurality of bases from distinct data sources." The cubes are defined as "multi-dimensional data structures that have elements comprising one or more aggregated dimensions and that are processed in a multi-dimensional database." The Examiner contends that Megiddo's "synthetic databases" teach Appellants' claimed cubes. Megiddo's synthetic databases, however, are defined as "subsets of the original database." Megiddo at col. 7 lines 13-15. Thus, Megiddo's synthetic databases do not represent multi-dimensional data structures that comprise a plurality of bases from distinct data sources, as required by claim 1. Castelli is similarly deficient.

Further, neither Megiddo nor Castelli teaches generating **scoped** association cubes, as required by claim 1. Megiddo teaches association rules that are not scoped, being defined by a "similar probability threshold value." See Megiddo at col. 3, lines 41-51. The Castelli reference is directed towards approximating data using multiple projections and/or representations and does not address association rules. For either or both of these reasons, the Examiner erred in rejecting claim 1 and its dependent claims.

D. Claim 8

Claim 8 requires, among other features, a plurality of local stations and at least one global station coupled to the plurality of the local stations. The local stations have a local computation engine for mining and summarizing the local transaction data and for generating local customer profile cubes, and the global station has a global computation engine for receiving the local customer profiles, merging and mining the local profile cubes, and generating global profile cubes and scoped association rules. The scoped association rules comprise a plurality of bases from distinct data sources and are based on said local profile cubes. As articulated above regarding claim 1, neither Megiddo nor Castelli teach or even suggest generating scoped association rules, as required by claim 8.

Further, neither Megiddo nor Castelli teach or even suggest distributing the processing to a local and global station, as required by claim 8. Thus, claim 8 possesses structural and functional limitations not found in the art of record. At least for these reasons, the Examiner erred in rejecting claim 8 and its dependent claims.

E. Claim 17

Claim 17 requires, among other features, mining and summarizing, using a plurality of local servers, and merging and mining, using at least one global server. In addition, the global sever generates scoped association rules. As discussed above regarding claim 8, neither Megiddo nor Castelli teach or suggest distributing the processing to local and global servers. Further, the art of record does not teach or suggest the generation of scoped association rules, as required by claim 17. At least for these reasons, the Examiner erred in rejecting claim 17 and its dependent claims.

F. Claim 28

Claim 28 requires, among other features, a plurality of local stations having a local computation engine coupled to at least one global station having a global computation engine. The global computation engine generates scoped association rules. As discussed above regarding claim 8, neither Megiddo nor Castelli teach or suggest distributing processing to local and global stations. Further, neither Megiddo nor Castelli teaches the generation of scoped association rules, as required by claim 28. At least for these reasons, the Examiner erred in rejecting claim 28 and its dependent claims.

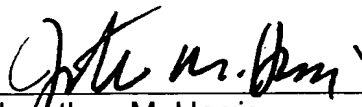
IX. CONCLUSION

For the reasons stated above, Appellants respectfully submit that the Examiner erred in rejecting all pending claims. If any fees or time extensions are inadvertently omitted or if any fees have been overpaid, please appropriately charge or credit those fees to Hewlett-Packard Company Deposit Account

Appl. No.: 09/524,140
Appeal Brief dated June 11, 2004
Reply to Office action of February 19, 2004

Number 08-2025 and enter any time extension(s) necessary to prevent this case from being abandoned.

Respectfully submitted,


Jonathan M. Harris
PTO Reg. No. 44,144
CONLEY ROSE, P.C.
(713) 238-8000 (Phone)
(713) 238-8008 (Fax)
ATTORNEY FOR APPELLANTS

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
Legal Dept., M/S 35
P.O. Box 272400
Fort Collins, CO 80527-2400



APPENDIX TO APPEAL BRIEF
CURRENT CLAIMS

1. (Previously presented) A method for generating association rules comprising:

in a processing system, receiving a volume cube that represents the purchase volume of customers;

in the processing system, generating scoped association cubes, a population cube and a base cube based on the volume cube, wherein the scoped association cubes comprise a plurality of bases from distinct data sources; and

in the processing system, deriving a confidence cube and a support cube of an association rule based on the association cube, population cube, and the base cube,

wherein said volume cube, association cube, population cube, base cube, and confidence cube comprise multi-dimensional data structures that have elements comprising one or more aggregated dimensions and that are processed in a multi-dimensional database.

2. (Original) The method of claim 1 wherein generating an association cube, a population cube and a base cube based on the volume cube comprises generating an association cube that has at least two levels and at least two dimensions.

3. (Original) The method of claim 1 wherein generating an association cube, a population cube and a base cube based on the volume cube comprises generating a scoped association rule cube and wherein deriving a confidence cube and a support cube of an association rule based on the association cube, population cube, and the base cube comprises deriving a confidence cube and a support cube of a scoped association rule based on the association cube, population cube, and the base cube.

4. (Original) The method of claim 1 wherein generating an association cube, a population cube and a base cube based on the volume cube comprises:
 - generating an association rule with conjoint items cube;
 - wherein deriving a confidence cube and a support cube of an association rule based on the association cube, population cube, and the base cube comprises:
 - deriving a confidence cube and a support cube of an association rule with conjoint items based on the association cube, population cube, and the base cube.
5. (Original) The method of claim 1 wherein generating an association cube, a population cube and a base cube based on the volume cube comprises:
 - generating a functional association rule cube;
 - wherein deriving a confidence cube and a support cube of an association rule based on the association cube, population cube, and the base cube comprises:
 - deriving a confidence cube and a support cube of a functional association rule based on the association cube, population cube, and the base cube.
6. (Original) The method of claim 1 wherein steps receiving, generating and deriving are implemented by utilizing a OLAP programming.
7. (Original) The method of claim 1 wherein receiving comprises:
 - receiving a first volume cube that represents the purchase volume of customers for a first region;
 - receiving a second volume cube that represents the purchase volume of customers for a second region; and
 - wherein generating comprises:
 - generating an association cube, a population cube and a base cube based on the first volume cube and the second volume cube.

8. (Previously presented) A data processing system comprising:
a plurality of local stations ("LDOSs") having a local computation engine for mining and summarizing the local transaction data and for generating local customer profile cubes; and
at least one global station ("GDOS"), coupled to the plurality of the local stations, the global station having a global computation engine for receiving the local customer profiles, merging and mining the local profile cubes, and generating global profile cubes and scoped association rules, the scoped association rules comprising a plurality of bases from distinct data sources and based on said local profile cubes, and providing the global profile cubes and the association rules to said plurality of LDOSs,
wherein said local customer profile cubes and global profile cubes comprise multi-dimensional data structures that have elements comprising one or more aggregated dimensions and that are processed in a multi-dimensional database.
9. (Original) The system according to claim 8, wherein each of said plurality of LDOSs comprises a local data warehouse and at least one local OLAP server, the local data warehouse being adapted to receive and store said transaction data,
wherein the local computation engine builds the local profile cubes that contains at least partial information regarding customer profiling by periodically mining new transactions flowing into said local data warehouse and deriving patterns for local analysis, said local computation engine also being adapted to incrementally update said local profile cubes.
10. (Original) The system according to claim 9 wherein said local data warehouse receives and stores transaction data in a first predetermined interval

and wherein said local OLAP engine generates said local profile cubes in a second predetermined interval.

11. (Original) The system according to claim 9 wherein said GDOS comprises a global data warehouse and at least one global OLAP server,

the global data warehouse for receiving and storing the local profile cubes,
the global computation engine for combining summary information from each of said LDOSs to build and incrementally update said global profile cubes and association rules, and for providing feedback to said plurality of LDOSs.

12. (Original) The system according to claim 11, wherein said local and global profile cubes comprise information of a plurality of customers, said information being derived from transaction data with said customers as stored by said local and global data warehouses, said profiling information specifying at least the following: kind, product, customer, merchant, time and area.

13. (Original) The system according to claim 12, wherein:

said local profile cubes are maintained at LDOS and said global profile cubes are maintained at GDOS, each of said local profile cubes being populated by mapping values in transaction data records into each dimension of said profile cube, each of said global profile cubes being retrieved and updated by merging appropriate local profile cubes.

14. (Original) The system according to claim 12, wherein said profile cubes are used to derive a plurality of shopping pattern cubes, said shopping pattern cubes comprising:

shopping behavior of at least one customer;
shopping patterns based on probability distribution;
shopping patterns based on volume.

15. (Original) The system according to claim 8, wherein said association rules comprise:

- scoped association rule with different bases, each of the bases being said scoped association rule's population over which said scoped association rule is defined;

- multidimensional association rule with "customer" being its base, "products" being its item, and "merchant," "area" and "time" being underlying features of said multidimensional association rule; and

- multilevel association rule with its features being represented at multiple levels.

16. (Original) The system according to claim 15, wherein said association rules are mined by:

- converting a volume cube into an association cube, a base cube and a population cube, said volume cube representing purchase volumes of customers dimensioned by item, base and feature;

- deriving a support cube based on said base cube and said association cube; and

- deriving a confidence cube based on said association cube and said population cube.

17. (Previously presented) A method of distributed data processing using on-line analytical processing ("OLAP") engines for use with transaction data in electronic commerce, comprising the steps of:

- mining and summarizing, using a plurality of local servers ("LDOSs"), said transaction data to generate local profile cubes;

- merging and mining, using at least one global server ("GDOS"), said local profile cubes received from said plurality of LDOSs to generate global profile cubes and scoped association rules based on said local profile cubes, wherein the scoped association rules comprise a plurality of bases from distinct data sources; and

feeding back said global profile cubes and association rules from said GDOS to said plurality of LDOSs for their business applications, wherein said local profile cubes and global profile cubes comprise multi-dimensional data structures that have elements comprising one or more aggregated dimensions and that are processed in a multi-dimensional database.

18. (Original) The method according to claim 17, wherein the step of mining and summarizing, using LDOSs, comprises:

receiving and storing said transaction data using a local data warehouse, building, using a local OLAP engine, said local profile cubes containing at least partial information regarding customer profiling by periodically mining new transactions flowing into said local data warehouse and deriving patterns for local analysis; and incrementally updating said local profile cubes with the new transactions.

19. (Original) The method according to claim 18 wherein the step of receiving and storing is in a first predetermined interval and wherein the step of building said local profiles is in a second predetermined interval.

20. (Original) The method according to claim 18 wherein said step of receiving, merging and mining by using the GDOS comprises:

storing said local profile cubes using a global data warehouse; combining using a global OLAP engine summary information from each of said LDOSs to build and incrementally update said global profile cubes and association rules; and feeding back said global profile cube and association rules.

21. (Original) The method according to claim 20 wherein said local and global profile cubes include information of a plurality of customers, said information being derived from transaction data with said customers as stored by said local

and global data warehouses, said profiling information specifying at least the following: kind, product, customer, merchant, time and area.

22. (Original) The method according to claim 21 wherein said local profile cubes are maintained at LDOS and said global profile cubes are maintained at GDOS, each of said local profile cubes being populated by mapping values in transaction data records into each dimension of said profile cube, each of said global profile cubes being retrieved and updated by merging appropriate local profile cubes.

23. (Original) The method according to claim 21 wherein said profile cubes are used to derive a plurality of shopping pattern cubes, said shopping pattern cubes comprising:

- shopping behavior of at least one customer;
- shopping patterns based on probability distribution;
- shopping patterns based on volume.

24. (Original) The method according to claim 17 wherein said association rules comprise:

- scoped associate rule with different bases, each of the bases being said scoped association rule's population over which said scoped association rule is defined;
- multidimensional association rule with "customer" being its base, "products" being its item, and "merchant," "area" and "time" being underlying features of said multidimensional association rule; and
- multilevel association rule with its features being represented at multiple levels.

25. (Original) The method according to claim 24, wherein said association rules are mined by:

converting a volume cube into an association cube, a base cube and a population cube, said volume cube representing purchase volumes of customers dimensioned by item, base and feature;
deriving a support cube based on said base cube and said association cube; and
deriving a confidence cube based on said association cube and said population cube.

26. (Previously presented) A method for generating association rules, comprising:

in the processing system, generating scoped association cubes, a population cube and a base cube based on a volume cube that represents the purchase volume of customers, wherein the scoped association cubes comprise a plurality of bases from distinct data sources; and
in the processing system, deriving a confidence cube and a support cube of an association rule based on the association cube, population cube, and the base cube.

27. (Previously presented) The method of claim 26 wherein generating scoped association cubes, a population cube, and a base cube based on a volume cube comprises generating an association cube that has at least two levels and at least two dimensions.

28. (Previously presented) A system, comprising:

a plurality of local stations having a local computation engine for mining and summarizing the local transaction data and for generating local customer profile cubes; and
at least one global station coupled to the plurality of the local stations, the global station having a global computation engine for receiving the local customer profiles, merging and mining the local profile cubes,

and generating global profile cubes and scoped association rules, the scoped association rules comprising a plurality of bases from distinct data sources and based on said local profile cubes.

29. (Previously presented) The system of claim 29 wherein the local and global stations comprise computer servers adapted to perform aggregations on the local customer profile cubes.